

10-11-00

A

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
P4272; 03226/047001

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTSBox Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

IMPROVING THE ACCURACY OF MULTIPLE BRANCH PREDICTION SCHEMES

and invented by:

Adam R. TALCOTTIf a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 21 pages and including the following:
- a. ☒ Descriptive Title of the Invention
 - b. ☐ Cross References to Related Applications (if applicable)
 - c. ☐ Statement Regarding Federally-sponsored Research/Development (if applicable)
 - d. ☐ Reference to Microfiche Appendix (if applicable)
 - e. ☒ Background of the Invention
 - f. ☒ Brief Summary of the Invention
 - g. ☒ Brief Description of the Drawings (if drawings filed)
 - h. ☒ Detailed Description
 - i. ☒ Claim(s) as Classified Below
 - j. ☒ Abstract of the Disclosure

**22511**

PATENT TRADEMARK OFFICE

UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
P4272; 03226/047001

Total Pages in this Submission

Application Elements (Continued)

3. ☒ Drawing(s) (when necessary as prescribed by 35 USC 113)

- a. ☒ Formal Number of Sheets 5
- b. ☐ Informal Number of Sheets _____

4. ☒ Oath or Declaration

- a. ☒ Newly executed (original or copy) ☐ Unexecuted
- b. ☐ Copy from a prior application (37 CFR 1.63(d)) (for continuation/divisional application only)
- c. ☐ With Power of Attorney ☒ Without Power of Attorney
- d. ☐ DELETION OF INVENTOR(S)

Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).

5. ☐ Incorporation By Reference (usable if Box 4b is checked)

The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.

6. ☐ Computer Program in Microfiche (Appendix)

7. ☐ Nucleotide and/or Amino Acid Sequence Submission (if applicable, all must be included)

- a. ☐ Paper Copy
- b. ☐ Computer Readable Copy (identical to computer copy)
- c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. ☒ Assignment Papers (cover sheet & document(s))

9. ☐ 37 CFR 3.73(B) Statement (when there is an assignee)

10. ☐ English Translation Document (if applicable)

11. ☐ Information Disclosure Statement/PTO-1449 ☐ Copies of IDS Citations

12. ☐ Preliminary Amendment

13. ☒ Acknowledgment postcard

14. ☒ Certificate of Mailing

☐ First Class ☒ Express Mail (Specify Label No.): EL521607231US

**UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)**

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
P4272; 03226/047001

Total Pages in this Submission

Accompanying Application Parts (Continued)

15. ☐ Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. ☐ Additional Enclosures *(please identify below):*

Request That Application Not Be Published Pursuant To 35 U.S.C. 122(b)(2)

17. ☐ Pursuant to 35 U.S.C. 122(b)(2), Applicant hereby requests that this patent application not be published pursuant to 35 U.S.C. 122(b)(1). Applicant hereby certifies that the invention disclosed in this application has not and will not be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication of applications 18 months after filing of the application.

Warning

An applicant who makes a request not to publish, but who subsequently files in a foreign country or under a multilateral international agreement specified in 35 U.S.C. 122(b)(2)(B)(i), must notify the Director of such filing not later than 45 days after the date of the filing of such foreign or international application. A failure of the applicant to provide such notice within the prescribed period shall result in the application being regarded as abandoned, unless it is shown to the satisfaction of the Director that the delay in submitting the notice was unintentional.

UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
P4272; 03226/047001

Total Pages in this Submission


Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	33	- 20 =	13	x \$18.00	\$234.00
Indep. Claims	6	- 3 =	3	x \$80.00	\$240.00
Multiple Dependent Claims (check if applicable) <input type="checkbox"/>					\$0.00
BASIC FEE					\$710.00
OTHER FEE (specify purpose) <u>Assignment Fee</u>					\$40.00
TOTAL FILING FEE					\$1,224.00

- ☒ A check in the amount of **\$1,224.00** to cover the filing fee is enclosed.
- ☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. **50-0591** as described below. A duplicate copy of this sheet is enclosed.
- ☐ Charge the amount of _____ as filing fee.
- ☒ Credit any overpayment.
- ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

Dated: 10/9/00


Signature 45,079

Jonathan P. Osha, Reg. No. 33,986
ROSENTHAL & OSHA L.L.P.
700 Louisiana, Suite 4550
Houston, Texas 77002

Telephone: (713) 228-8600
Facsimile: (713) 228-8778

CC:

CERTIFICATE OF MAILING BY "EXPRESS MAIL" (37 CFR 1.10)

Applicant(s): Adam R. TALCOTT

Docket No.

P4272; 03226/047001

Serial No.

Filing Date

Examiner

Group Art Unit

Invention: **IMPROVING THE ACCURACY OF MULTIPLE BRANCH PREDICTION SCHEMES**


I hereby certify that the following correspondence:

Utility Patent Application

(Identify type of correspondence)

is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under

37 CFR 1.10 in an envelope addressed to: The Assistant Commissioner for Patents, Washington, D.C. 20231 on

October 9, 2000*(Date)*Lisa H. Smith*(Typed or Printed Name of Person Mailing Correspondence)**(Signature of Person Mailing Correspondence)*EL521607231US*("Express Mail" Mailing Label Number)***Note: Each paper must have its own certificate of mailing.**

**APPLICATION
FOR
UNITED STATES LETTERS PATENT**

**TITLE: IMPROVING THE ACCURACY OF MULTIPLE BRANCH
PREDICTION SCHEMES**

APPLICANTS: Adam R. TALCOTT



22511

PATENT TRADEMARK OFFICE

"EXPRESS MAIL" Mailing Label Number: EL521607231US
Date of Deposit: October 9, 2000

IMPROVING THE ACCURACY OF MULTIPLE BRANCH PREDICTION SCHEMES

BACKGROUND OF THE INVENTION

Computer processors contain arithmetic, logic, and control circuitry that interpret and execute instructions from a computer program. Referring to Figure 1, a typical computer system includes a microprocessor (22) having, among other things, a CPU (24), a memory controller (26), and an on-board cache memory (30). The microprocessor (22) is connected to external cache memory (32) and a main memory (34) that both hold data and program instructions to be executed by the microprocessor (22). Internally, the execution of program instructions is carried out by the CPU (24). Data needed by the CPU (24) to carry out an instruction are fetched by the memory controller (26) and loaded into internal registers (28) of the CPU (24). Upon command from the CPU (24), the memory controller (26) searches for the data first in the fast on-board cache memory (30), then in external cache memory (32), and finally in the slow main memory (34). Finding the data in the cache memory is referred to as a "hit." Not finding the data in the cache memory is referred to as a "miss."

The time between when a CPU requests data and when the data is retrieved and available for use by the CPU is termed the "latency" of the system. If requested data is found in cache memory, *i.e.*, a data hit occurs, the requested data can be accessed at the speed of the cache and the latency of the system is reduced. If, on the other hand, the data is not found in cache, *i.e.*, a data miss occurs, and thus the data must be retrieved from main memory for access and the latency of the system is increased.

In the pursuit of improving processor performance, designers have sought two main goals: making operations faster and executing more operations in parallel. Making operations faster can be approached in several ways. For example, transistors can be made to switch faster and thus propagate signals faster by improving semiconductor processes; execution-unit latency can be reduced by increasing the number of transistors in the design; and the levels of logic required by the design to implement a given function can be minimized to increase speed. To execute more operations in parallel, designers mainly rely on one, or a combination of pipelining and superscalar techniques. Pipelined processors overlap instructions in time on common execution resources. Superscalar processors overlap instructions in space on separate resources.

Pipeline stalls are a main performance inhibitor with regard to parallel processing. Stalls arise from data dependencies, changes in program flow, and hardware resource conflicts. At times, pipeline stalls can be avoided by rearranging the order of execution for a set of instructions. Compilers can be used to statically reschedule instructions. However, incomplete knowledge of run-time information reduces the effectiveness of static rescheduling. In-order processors, *i.e.*, processors that issue, execute, complete, and retire instructions in strict program order, have to rely entirely on static rescheduling and thus are prone to pipeline stalls.

As a result, designers use out-of-order processors and seek to implement dynamic instruction rescheduling. The simplest out-of-order processors issue instructions in order but allow them to execute and complete out of order. Even these simple out-of-order processors require complex hardware to reorder results before the corresponding instructions are retired. A strict result order is not required from a data-flow perspective.

However, such ordering is necessary to maintain precise exceptions and to recover from mispredicted speculative execution.

A well-known method of reordering is through the use of a reorder buffer, *i.e.*, a buffer that maintains results until written to the register file in program order. Designers also use other types of reordering hardware, such as history buffers and future files. History buffers record source-operand history so the processor can backtrack to a precise architectural state and future files store the current state and the architectural state in separate register files allowing the processor to be restored to a precise check-point state.

Branch prediction and speculative execution are additional techniques used to reduce pipeline stalls. In a pipelined processor, the outcomes of branch instructions are often determined after subsequent instructions have been fetched. Using branch prediction schemes, microprocessors attempt to accurately predict whether a branch is taken or not based on how that branch has behaved previously. The aggregate behavior, or the average behavior over time, of the branch instruction is stored in a Branch Prediction Table ("BPT"). Given a branch instruction's aggregate behavior, the branch predictor, which resides in an instruction fetch unit, predicts the outcome of the branch instruction and then loads instructions thereafter based on that prediction. For example, if the branch predictor predicts that a branch will be taken, then the processor fetches subsequent instructions according to the address to which the instruction branches. When the branch proceeds in the predicted direction, pipeline stalls are completely avoided. On the other hand, if the branch direction is mispredicted, all the instructions after the mispredicted instruction must be removed from the processor.

Modern microprocessors incorporate a variety of branch prediction schemes.

These schemes usually fall under one of two broad classifications: static branch prediction and dynamic branch prediction. Static branch prediction occurs when a branch predictor makes predictions that are not based on the run-time behavior of branch instructions. Two such schemes are: “predict not taken” and “predict taken.” In the “predict not taken” scheme, a branch is predicted as not taken, and the processor simply continues as if the branch did not exist. In the “predict taken” scheme, as soon as the branch is decoded and the target address of the next instruction is predicted, it is assumed that the branch is taken and the process continues with the fetching and executing of instructions at the target address.

Dynamic branch prediction, on the other hand, occurs when the processor responds to changes in a branch instruction’s behavior while a program is executing. In other words, a dynamic branch prediction scheme provides a mechanism by which a processor can take into account the cumulative behavior of a branch instruction. In cases where there are more than one branch instruction in an instruction bundle, also known as a fetch bundle, some schemes may break the fetch bundle at the points where additional branch instructions reside.

Typically, dynamic branch prediction schemes are extended to predict multiple branches. The increasing number of instructions executed per cycle in high-performance microprocessors calls for the instruction fetch unit to fetch a correspondingly larger number of instructions each cycle. Fetching more instructions per cycle increases the likelihood that more than one branch prediction will need to be made each cycle to determine the next fetch address. Accordingly, a microprocessor handles multiple branch instructions through a multiple branch prediction scheme.

Multiple branch prediction schemes usually depend upon a fetch bundle address (FBA), *i.e.*, the address of the first instruction fetched in a given cycle, to index all prediction structures for the fetched instruction bundle. Each branch to be predicted in the current fetch bundle uses information from the prediction structures indexed by the address of the bundle currently being fetched. In other words, an instruction uses its fetch bundle address to base its index address, which it uses to index a branch prediction structure.

An index based upon the fetch bundle address has typically been used to select several prediction counters from which multiple branch predictions are made dependent upon the location of the branches in the fetch bundle. These counters may be simple two-bit counters with saturation. Typically, the counter is incremented when the branch instruction is taken and it is decremented when the instruction is not taken. For example, the branch instruction could be taken if the counter is in state 2 or 3, and not taken if the counter is in state 0 or 1. The saturation attribute inhibits the counter from recycling. For example, if the counter is in state 3 and a branch instruction is taken, the counter will stay in state 3 and not increment back to state 0. This mechanism attempts to predict branch instructions by gathering and representing the aggregate behavior of previous branch instructions. Improvements made to this mechanism include using branch history so that the processor recognizes repetitive patterns in previous branch instructions to predict later branch instructions.

Typically, branch history mechanisms are separated into two types: global branch history and local branch history. Global branch history pertains to the last g outcomes of

executed branch instructions. Local branch history incorporates the last l outcomes of the current branch instruction.

Branch instructions can either be conditional or unconditional. All conditional branch instructions and most unconditional branch instructions can be incorporated into branch history registers. However, some unconditional branch instructions, such as call and return statements, are usually not incorporated into branch history registers. Global branch history prediction schemes have been extended to predict the outcomes of multiple branch instructions. Additionally, in order to predict multiple conditional branch instructions simultaneously, the BPT is typically extended to contain several n -bit counters per entry. Depending on the location of branches in the current fetch bundle, the proper n -bit counters are referenced to make the multiple branch predictions.

Figure 2 depicts instruction addresses A through A+9, which constitute an exemplary instruction fetch bundle. An instruction address is the location of where an instruction resides in memory. The two exemplary instruction fetch blocks are indexed A (36) and A+5 (40). If the branch instruction resides at A+4 (38), and the branch is predicted taken and mispredicted, the instruction at A+6 (42) uses A+5 (40) as its index address. Alternatively, if a branch instruction is mispredicted, then a subsequent branch instruction depends on a new fetch bundle address to base its index address. Therefore, the prediction information associated with the subsequent branch instruction is different than the prediction information it would have used if it used a prior fetch bundle address. This is due to the fact that the subsequent branch instruction uses prediction information that is based on a newly indexed branch prediction table. One skilled in the art will

appreciate that the described method is illustrative of one typical indexing method and, depending on the processor architecture, other methods can be used to index instructions.

SUMMARY OF THE INVENTION

5 In one aspect, a method of improving the prediction accuracy of a branch prediction scheme comprises reading an individual instruction in a current set of instructions, fetching the individual instruction when an instruction fetch unit determines that the individual instruction is valid, and allowing the instruction fetch unit to use an index address for the fetched individual instruction.

10 In one aspect, a method of improving branch prediction accuracy comprises receiving a set of instructions having an assigned address, making a prediction for a branch instruction in the set of instructions using the assigned address, and retaining the assigned address for the branch instruction in the set of instructions.

15 In one aspect, a tool for improving prediction accuracy of a branch prediction scheme comprises a processor for reading an individual instruction in a current set of instructions, and an instruction fetch unit for determining whether the individual instruction is valid and fetching the individual instruction when the individual instruction is valid, wherein an index address is used for the fetched individual instruction.

20 In one aspect, a tool of improving branch prediction accuracy comprises a set of instructions having an assigned address, and a branch predictor for making a prediction for a branch instruction in the set of instructions using the assigned address, wherein the assigned address for the branch instruction in the set of instructions is retained.

In one aspect, an apparatus for improving prediction accuracy of a branch instruction scheme comprises means for reading an individual instruction in a current set of instructions, means for fetching the individual instruction when an instruction fetch unit determines that the individual instruction is valid, and means for allowing the instruction fetch unit to use an index address for the fetched individual instruction.

In one aspect, an apparatus for improving branch prediction accuracy comprises means for receiving a set of instructions having an assigned address, means for predicting an outcome for a branch instruction in the set of instructions using the assigned address, and means for retaining the assigned address for the branch instruction in the set of instructions.

Other advantages and features will become apparent from the following description, including the figures and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a typical computer system.

Figure 2 is a block diagram of an exemplary instruction bundle.

Figure 3 shows a computer system in accordance with an embodiment of the present invention.

Figure 4 is a flow chart describing a branch instruction indexing process in accordance with the present invention.

Figure 5 is a block diagram of an exemplary instruction bundle.

DETAILED DESCRIPTION

The present invention relates to a system for improving the accuracy of multiple branch prediction schemes. Referring to the drawings wherein like reference characters are used for like parts throughout the several views, Figure 3 is a computer system in accordance with an embodiment of the present invention. The system encourages the correct prediction behavior from branch prediction structures, which comprise prediction information, by maintaining the proper index components for all branch prediction structures even in the presence of mispredictions. In one aspect, the accuracy of a branch prediction scheme is increased by using the same prediction information to predict a branch instruction subsequent to a mispredicted branch instruction that would have been used if the mispredicted instruction had been correctly predicted.

Referring to Figure 3, a system in accordance with an embodiment of the present invention includes a microprocessor (46) having an instruction pipeline (56), and an instruction fetch unit (52) having a branch prediction table ("BPT") (48) and an instruction cache (49). Also, the system interfaces with memory units (44) similar to the microprocessor described in Figure 1. The instruction pipeline (56) consists of a series of Decode units (D1...Dn) and Execution units (E1...En).

As instructions become available to the pipeline (56), the pipeline (56) decodes and executes the instructions. The instruction fetch unit (52) reads out history information when an instruction bundle needs to be fetched. Typically, an index based upon a fetch bundle address is used to select several prediction counters in the BPT (48) from which multiple predictions are made based upon the location of branches in a fetch bundle.

Figure 4 is a flow diagram of an exemplary branch instruction indexing process in accordance with an embodiment of the present invention. Initially, as an instruction fetch unit begins to read individual instructions sequentially from an instruction cache for a given cycle (step 70), the instruction fetch unit decodes each individual instruction (step 72). Those skilled in the art will appreciate that in other embodiments, the instruction fetch unit instead of decoding an instruction may use pre-decode information, *i.e.*, decode information that is computed as instructions are written into the instruction cache in the instruction fetch unit and before the instruction reaches branch prediction logic. As instructions are individually read and validated by the instruction fetch unit, the instructions are subsequently said to be processed, *i.e.*, fetched, and form an instruction fetch bundle. Validating an instruction comprises, among other things, making sure the instruction does not follow a predicted taken branch. Associated with the instruction fetch bundle is a fetch bundle address. The fetch bundle address is the address of the first instruction fetched for the given cycle. If the decode information for a particular instruction reveals that the instruction is not a branch instruction (step 74), the instruction fetch unit processes the instruction and proceeds with reading and decoding the next instruction in the instruction cache (step 80). However, if the decode information indicates that the instruction is a branch instruction (step 74), the instruction fetch unit must make a prediction as to whether or not the branch instruction branches. To make the prediction, the instruction fetch unit must determine a proper address to use to index a branch prediction structure.

In determining the proper index address for the branch instruction, the instruction fetch unit must determine whether the branch instruction was read before when a prior

instruction fetch bundle (step 76) was formed. If the branch instruction was not read in a previous cycle (step 76), the instruction fetch unit processes the instruction, wherein the branch instruction becomes part of an instruction fetch bundle, and then uses the fetch bundle address of that instruction fetch bundle (step 79) to index a branch prediction structure. However, if the instruction fetch unit read the branch instruction in a prior cycle (step 76), the instruction fetch unit makes a further determination as to the proper fetch bundle address the branch instruction should use (step 78) when indexing a branch prediction structure to make a prediction for the branch instruction. In other words, the instruction fetch unit determines what the fetch bundle address was when the branch instruction was first read.

For example, when an instruction fetch unit reads and processes a set of 8 instructions (hereafter referred to as the “original” instruction set), and where the second (the “first” branch instruction) and fifth (the “next” branch instruction) instructions in the original instruction set are branch instructions, if the first branch instruction is predicted taken and mispredicted, the next branch instruction is removed from the instruction fetch unit. That is, because of the misprediction, all instructions that were read after the mispredicted instruction are removed from the instruction fetch unit, and consequently, the instruction fetch unit does not process the removed instructions. Thereafter, when the next branch instruction in the original instruction set is re-fetched as part of another instruction fetch bundle, the instruction fetch unit determines that the re-fetched branch instruction should use the fetch bundle address that would have been used if the first branch instruction were predicted correctly. The process is explained in further detail below.

Those skilled in the art will appreciate that the exemplary flow process shown in Figure 4 is equally applicable to a plurality of instructions being processed in parallel. Further, while, for purposes of illustration, steps are depicted occurring in series, those skilled in the art will appreciate that, according to dependency, certain steps may occur simultaneously.

Figure 5 represents an instruction sequence from addresses B through B+9, with a corresponding fetch bundle address of B (82). In a system in accordance with an embodiment of the present invention, if a misprediction occurs for the branch instruction at B+2 (84), then the next branch instruction occurring at B+7 (88) uses B (82) as its index. One skilled in the art will appreciate that although the initial misprediction causes the instruction fetch unit to fetch a new bundle of instructions with a new fetch bundle address, *i.e.*, B+3 (86), the instruction that initially occurred at B+7 (88) still uses B (82) as its index.

Accordingly, the second branch instruction at B+7 (88) uses the initial index address of B (82), and therefore, the branch predictor becomes more quickly trained because the branch predictor uses the prediction information that would have been used if an accurate prediction was made. That is, the counters associated with the branch instruction at B+7 (88) more quickly reflect the true behavior of the instruction without the deviation that might have occurred had the branch instruction at B+7 (88) used the B+3 (86) index because when B+2 (84) is later correctly predicted, B+7 (88) uses B (82) again. Those skilled in the art will appreciate that because the branch predictor becomes more quickly trained, the chances of mispredictions decrease. This is due to the fact that once a particular branch is correctly predicted, the fetch bundle is not broken and a

subsequent branch uses a portion of the prediction information that was used the previous times the branch was first encountered.

Further, by training a branch predictor to behave in the manner in which the predictor would have behaved if a preceding branch instruction had not been mispredicted, the branch predictor is more quickly trained, and therefore the overall accuracy of the branch prediction scheme increases. Once the initial instruction bundle is fetched, all branch instructions contained within that fetch bundle will have the same index, *i.e.*, the proper fetch bundle address. When a branch instruction after a mispredicted branch instruction is removed due to a misprediction, the branch instruction uses the fetch bundle address that it would have used if the misprediction did not occur to base its index address when fetched again in a different instruction bundle with a new fetch bundle address.

Because the system allows instructions to use the fetch bundle address which would have been used if a prior misprediction had not occurred, the corresponding branch prediction table is manipulated as if the microprocessor did not make a misprediction. This allows a branch predictor to adhere to the natural behavior of the instruction and not a behavior that is altered due to mispredictions. Further, by having the branch predictor determine which fetch bundle address a branch instruction should use, branch instructions can use proper fetch bundle addresses and therefore, the corresponding indexes allow the microprocessor to refer to the originally indexed branch prediction structure entries for the instructions. Because of this, prediction counters associated with the instructions more quickly go to a state that more accurately represents their natural behavior. Those

skilled in the art will appreciate that in other embodiments, the prediction process applies to other branch predictors, such as address predictors.

Advantages of the present invention may include one or more of the following. In one or more embodiments of the present invention, a branch instruction is allowed to use an index address that it would have used if prior mispredictions did not occur. Therefore, a branch prediction table associated with a particular branch instruction more quickly reflects the instruction's behavior. Further, because the branch prediction table becomes more accurate quicker, fewer mispredictions will occur. As the amount of mispredictions a microprocessor makes is decreased, the performance of the system is increased. Those skilled in the art will appreciate that the present invention also may include other advantages and features.

While the present invention has been described with respect to a limited number of preferred embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. The appended claims are intended to cover all such modifications and variations which occur to one of ordinary skill in the art.

What is claimed is:

- 1 1. A method of improving prediction accuracy of a branch prediction scheme,
2 comprising:
3 reading an individual instruction in a current set of instructions;
4 fetching the individual instruction when an instruction fetch unit
5 determines that the individual instruction is valid; and
6 allowing the instruction fetch unit to use an index address for the fetched
7 individual instruction.
- 1 2. The method of claim 1, wherein the individual instruction is a branch instruction,
2 the method further comprising:
3 determining whether the branch instruction has been read in a previous set
4 of instructions.
- 1 3. The method of claim 2, further comprising:
2 selectively using a fetch bundle address for a plurality of fetched
3 individual instructions as the index address for the branch
4 instruction based on whether the branch instruction has been read
5 in the previous set of instructions.
- 1 4. The method of claim 2, further comprising:
2 determining a proper index address to use if the branch instruction has
3 been read in the previous set of instructions.
- 1 5. The method of claim 4, further comprising:
2 determining the proper index address by determining the fetch bundle
3 address the branch instruction would have used if a prior branch
4 instruction in the previous set of instructions had not been
5 mispredicted.

1 6. The method of claim 1, wherein the branch prediction scheme is for predicting an
2 outcome of a branch instruction.

1 7. The method of claim 1, wherein the index address is used to index an entry in a
2 branch prediction structure.

1 8. The method of claim 3, wherein the fetch bundle address is an address of a first
2 instruction in the plurality of fetched individual instructions.

1 9. The method of claim 1, wherein the plurality of fetched individual instructions is
2 an instruction fetch bundle.

1 10. The method of claim 1, further comprising:
2 using decode information for the individual instruction to determine
3 whether the individual instruction is a branch instruction.

1 11. The method of claim 1, further comprising:
2 using pre-decode information for the individual instruction to determine
3 whether the individual instruction is a branch instruction.

1 12. A method of improving branch prediction accuracy, comprising:
2 receiving a set of instructions having an assigned address;
3 making a prediction for a branch instruction in the set of instructions using
4 the assigned address; and
5 retaining the assigned address for the branch instruction in the set of
6 instructions.

1

- 1 13. The method of claim 12 further comprising:
2 making a prediction for an other branch instruction in the set of
3 instructions using the assigned address; and
4 retaining the assigned address for the other branch instruction in the set of
5 instructions.
- 1 14. The method of claim 12, wherein the assigned address is a fetch bundle address.
- 1 15. The method of claim 13, wherein the branch instruction is mispredicted, the
2 method further comprising:
3 removing the set of instructions having the assigned address;
4 receiving a second set of instructions having a second assigned address;
5 recognizing a branch instruction in the second set of instruction as a
6 branch instruction in the removed set of instructions; and
7 making a prediction for the recognized branch instruction using the
8 assigned address of the removed set of instructions.
- 1 16. A tool for improving prediction accuracy of a branch prediction scheme,
2 comprising:
3 a processor for reading an individual instruction in a current set of
4 instructions; and
5 an instruction fetch unit for determining whether the individual instruction
6 is valid and fetching the individual instruction when the individual
7 instruction is valid,
8 wherein an index address is used for the fetched individual instruction.
- 1 17. The tool of claim 16, wherein the individual instruction is a branch instruction,
2 and the instruction fetch unit is further for determining whether the branch
3 instruction has been read in a previous set of instructions.

1 18. The tool of claim 17, further comprising:
2 a fetch bundle address for a plurality of fetched individual instructions,
3 wherein the fetch bundle address is selectively used as the index address
4 for the branch instruction based on whether the branch instruction
5 has been read in the previous set of instructions.

1 19. The tool of claim 17, further comprising:
2 a proper index address is used if the branch instruction has been read in
3 the previous set of instructions.

1 20. The tool of claim 19, further comprising:
2 determining the proper index address by determining the fetch bundle
3 address the branch instruction would have used if a prior branch
4 instruction in the previous set of instructions had not been
5 mispredicted.

1 21. The tool of claim 16, wherein the branch prediction scheme is for predicting an
2 outcome of a branch instruction.

1 22. The tool of claim 16 further comprising:
2 an entry in a branch prediction structure indexed by the index address.

1 23. The tool of claim 18, wherein the fetch bundle address is an address of a first
2 instruction in the plurality of fetched individual instructions.

1 24. The tool of claim 16, wherein the plurality of fetched individual instructions is an
2 instruction fetch bundle.

1 25. The tool of claim 16, further comprising:
2 decode information for the individual instruction,
3 wherein the decode information is used to determine whether the
4 individual instruction is a branch instruction.

1 26. The tool of claim 16, further comprising:
2 pre-decode information for the individual instruction,
3 wherein the pre-decode information is used to determine whether the
4 individual instruction is a branch instruction.

1 27. A tool of improving branch prediction accuracy, comprising:
2 a set of instructions having an assigned address; and
3 a branch predictor for making a prediction for a branch instruction in the
4 set of instructions using the assigned address,
5 wherein the assigned address for the branch instruction in the set of
6 instructions is retained.

1 28. The tool of claim 27 wherein the assigned address is a fetch bundle address.

1 29. The tool of claim 27, wherein the branch predictor is further for making a
2 prediction for another branch instruction in the set of instructions using the
3 assigned address, wherein the assigned address for the other branch instruction in
4 the set of instructions is retained.

1 30. The tool of claim 29, wherein the branch instruction is mispredicted, the tool
2 further comprising:
3 a second set of instructions having a second assigned address, and
4 an instruction fetch unit for
5 removing the set of instructions having the assigned address; and
6 recognizes a branch instruction in the second set of instruction as a
7 branch instruction in the removed set of instructions,

8 wherein makes a prediction for the recognized branch instruction using the
9 assigned address of the removed set of instructions.

1 31. The tool of claim 27 wherein the prediction made is a prediction of an outcome of
2 the branch instruction.

1 32. An apparatus for improving prediction accuracy of a branch instruction scheme,
2 comprising:

3 means for reading an individual instruction in a current set of instructions;
4 means for fetching the individual instruction when an instruction fetch unit
5 determines that the individual instruction is valid; and
6 means for allowing the instruction fetch unit to use an index address for
7 the fetched individual instruction.

1 33. An apparatus for improving branch prediction accuracy, comprising:

2 means for receiving a set of instructions having an assigned address;
3 means for predicting an outcome for a branch instruction in the set of
4 instructions using the assigned address; and
5 means for retaining the assigned address for the branch instruction in the
6 set of instructions.

ABSTRACT

A method and apparatus of improving prediction accuracy of a branch instruction scheme includes reading an individual instruction in a current set of instructions, fetching the individual instruction when an instruction fetch unit determines that the individual instruction is valid, and allowing the instruction fetch unit to use an index address for the
5 fetched individual instruction. A method and apparatus of improving branch prediction accuracy includes receiving a set of instructions having an assigned address, making a prediction for a branch instruction in the set of instructions using the assigned address, and retaining the assigned address for the branch instruction in the set of instructions.

03226.047001.20000621.02.doc




Figure 1
(Prior Art)




Figure 2
(Prior Art)

FIG. 3 is a block diagram of a processor 46 in accordance with one embodiment of the present invention. The processor 46 includes a memory unit 44, an instruction fetch unit 52, and a plurality of processing units 56. The memory unit 44 is connected to the instruction fetch unit 52 via a bus 48. The instruction fetch unit 52 includes a branch prediction table 48 and an instruction cache 49. The processing units 56 include a plurality of decode units (D1, D2, ..., Dn) and a plurality of execute units (E1, E2, ..., En). The instruction fetch unit 52 is connected to the processing units 56 via a bus 50.




Figure 3

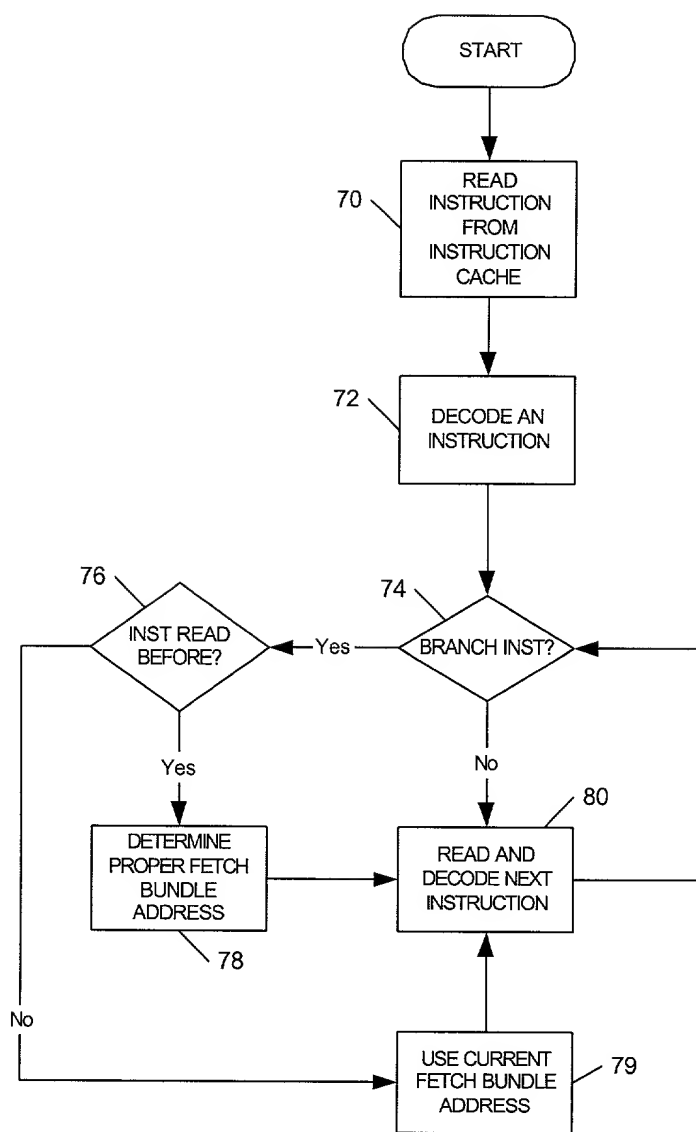


Figure 4




Figure 5

Docket No.

P4272/03226.047001

Declaration For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

EL521607231US

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for

IMPROVING THE ACCURACY OF MULTIPLE BRANCH PREDICTION SCHEMES



22511

PATENT TRADEMARK OFFICE

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on _____ as United States Application No. or PCT International Application Number _____ and was amended on _____

(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

(Number)

(Country)

(Day/Month/Year Filed)



(Number)

(Country)

(Day/Month/Year Filed)



(Number)

(Country)

(Day/Month/Year Filed)



I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, CFR Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or first inventor Adam R. Talcott	
Sole or first inventor's signature <i>Adam R. Talcott</i>	Date Oct 6, 2000
Residence 704 Villa Centre Way, San Jose, CA 95128	
Citizenshi USA	
Post Office 704 Villa Centre Way, San Jose, CA 95128	

Full name of second inventor, if any	
Second inventor's	Date
Residence	
Citizenshi	
Post Office	

Full name of third inventor, if any	
Third inventor's signature	Date
Residence	
Citizenshi	
Post Office	

Full name of fourth inventor, if any	
Fourth inventor's signature	Date
Residence	
Citizenshi	
Post Office	